

**ENERGY-AWARE WORKLOAD DISTRIBUTION****ABSTRACT OF THE DISCLOSURE**

5 The distribution of power dissipation within cluster systems is managed by a combination of inter-node and intra-node policies. The inter-node policy consists of subdividing the nodes within the cluster into three sets, namely the "Operational" set, the "Standby" set and the "Hibernating" set. Nodes in the Operational set continue to function and execute computation in response to user requests. Nodes in the Standby set have their processors in the low-energy standby mode and are ready to resume the computation immediately. Nodes in the Hibernating set are turned off to further conserve energy, and they need a relatively longer time to resume operation than nodes in the Standby set. The inter-node policy further distributes the computation among nodes in the Operational set such that each node in the set consumes the same amount of energy. Moreover, the inter-node policy responds to decreasing workload in the cluster by moving processors from the Operational set into the Standby set and by moving nodes from the Standby set into the Hibernating set. Vice versa, the inter-node policy responds to increasing workload in the cluster by moving nodes from the Hibernating set into the Operational set. Intra-node policies corresponding to managing the energy consumption within each node in the Operational nodes set by scaling operating frequency and power supply voltage corresponding to a given performance requirement.

10

15

::ODMA\PCDOCS\AUSTIN\_1\162662\3  
1097:7047/P424US